

UNITED STATES PATENT APPLICATION

of

Radia J. Perlman

and

Stephen R. Hanna

for a

**DATA AUTHENTICATION SYSTEM EMPLOYING ENCRYPTED INTEGRITY
BLOCKS**

DATA AUTHENTICATION SYSTEM EMPLOYING ENCRYPTED INTEGRITY BLOCKS

FIELD OF INVENTION

The invention relates generally to data communications networks and, in particular, to systems for authentication of data transferred over the networks.

BACKGROUND OF THE INVENTION

When a sender transfers data over a network, an interloper may intercept and alter the data and then transfer the altered data to the intended recipients. The recipients, who presume the data are valid because they appear to come from a trusted sender, may then use the altered data directly, or introduce errors into associated data processing systems. To ensure that the received data were not altered enroute by an interloper, the network may include a data authentication system that essentially encodes the data at the sender and decodes the data at the recipient to detect changes in the data.

One known data authentication process involves including a digital signature in a data packet. The digital signature is produced by first encoding the packet data bytes to produce a cryptographic hash and then, typically, encrypting the hash using the sender's private key. A recipient uses the sender's public key to decrypt the digital signature and reproduce the hash. It then encodes the received data using the same cryptographic hash function and compares the result with the decrypted hash. If the two hashes match, the data is considered authentic, that is, the received data is considered to be the same data that was sent by the sender who holds the private key. This authentication process works well, but it is both computation intensive and time consuming at the recipient end.

A slightly less computation intensive authentication process that may be used by senders and recipients that share a secret key is commonly referred to as a message integrity code, or "MIC," process. The MIC process produces an integrity code by concatenating the shared secret key with the data and then encoding the data and the key using a cryptographic hash function. The result, which is the integrity code, is then sent along with the data to a recipient who shares the secret key. The recipient similarly concatenates the shared secret key with the received data and encodes the data and the key using the hash function. If the result matches the received integrity code, the data is considered authentic. The MIC authentication process works well and is relatively reliable, assuming the holders of the shared key are trusted. However, this process still requires producing a hash based on all of the data bytes in the data packet, and is thus still relatively time consuming.

SUMMARY OF THE INVENTION

The invention is a data authentication system that at the sender produces for one or more data packets a plurality of "integrity checks" that the system then encrypts with a shared secret key, to produce an "integrity block." The system sends the integrity block as part of a data packet or separately. A recipient decrypts the integrity block using the shared secret key and reproduces the integrity checks. It then uses the integrity checks to authenticate the associated data packets. As discussed below, the authentication system uses a relatively weak, and thus, fast and uncomplicated, integrity check function to produce the integrity check, and adds robustness by selecting the integrity check function from a set of known functions and/or selecting from a data packet particular data bytes to manipulate using the selected function. Further robustness is added by encrypting the selection information into integrity checks, which ensures that the selections may vary between packets and that the particular selections are unknown to an interloper.

The recipient performs a single decryption operation and, based on the result, selects the appropriate integrity check function and/or data bytes for the one or more data packets. The recipient then performs one or more fast integrity check operations to authenticate the one or more associated data packets. Accordingly, the authentication

system operations are less time consuming and less complex than those of known prior authentication systems of similar robustness.

More specifically, the system selects an integrity check function from a set of known functions and a number of bytes from a given packet and manipulates the bytes in accordance with the selected integrity check function to produce the associated integrity check. The system next, as appropriate, selects from another packet corresponding bytes or bytes that are offset from the corresponding bytes, and produces a next associated integrity check using the same or another selected integrity check function, and so forth. The system then encrypts the integrity checks using a shared secret key, to produce the integrity block, which the system sends along with the one or more data packets to the intended recipients.

A recipient uses the shared secret key to decrypt the integrity block and reproduce the integrity checks. It then associates the integrity checks with the data packets and, for each data packet, selects the appropriate data bytes and authenticates the data using the appropriate integrity check function.

If an interloper intercepts the data packets and the integrity block, the interloper cannot readily interpret the integrity block because of the encryption. Accordingly, the interloper cannot readily determine, for example, which integrity check function to use, which bytes to select for authentication, and thus, how to alter the data and still pass the integrity checks without knowing the secret key. The unknown aspects of the authentication system inhibit the actions of the interloper, even if a relatively simple integrity check function is used and/or only a relatively small number of bytes from a given data packet are used, to produce the associated integrity check.

As discussed in more detail below, the integrity block may include information that is used by the recipient to determine, for a given data packet, which of the integrity check functions was selected and/or which of the data bytes were used to produce the associated integrity check. As necessary, the integrity block includes offset values, interval values and so forth. In addition, the integrity block may also include information, such as packet sequence numbers, that associate the integrity checks with the appropriate data packets. As discussed, all of the information in the integrity block, including the function

identifier, the offset or interval values and the sequence numbers, as appropriate, are preferably encrypted and, therefore, unknown to the interloper.

As a further safeguard against an interloper, the authentication system may include in the integrity block a digital signature that is prepared using the sender's private key. The recipient can then test that the integrity block has not been altered by an interloper that knows the shared secret key. Further, the sender may include digital signatures in the data packets so that the recipient can more robustly test selected data packets.

Alternatively, or in addition, the sender may include in a transmission one or more extraneous, or "chaff," data packets, which are data packets that intentionally fail the associated integrity checks. The sender may, for example, include in a transmission multiple sets of packets with the same sequence numbers. The recipient readily determines which of the packets with the same sequence numbers are valid using the appropriate integrity check. However, an interloper who cannot decipher the integrity block cannot as easily determine which of the packets are valid, and thus, cannot determine which packets to alter and/or how to alter these packets without detection.

To avoid having an interloper reproduce a valid data packet and flood the network with copies, a time stamp may be encrypted into the integrity check. The receiving station then, as discussed in more detail below, checks that the time stamp is valid before passing on or processing the associated data packet or packets. Alternatively, the system may use packet sequence numbers for this purpose, as well as, for associating the integrity checks with the appropriate data packets.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention description below refers to the accompanying drawings, of which:

Fig. 1 is a functional block diagram of a network constructed in accordance with the invention;

Fig. 2 is a more detailed block diagram of an authentication system that is included in the network of Fig. 1;

Fig. 3 depicts a data packet that includes an integrity block produced by the system of Fig. 2; and

Fig. 4 is a functional block diagram of an alternative authentication system.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

5 Referring to Fig. 1, a communications system 10 includes a plurality of end stations 12 that act as either senders or recipients or both. A given sending end station 12 sends data, in the form of data packets, through a communications network 16 to one or more selected receiving end stations 12 that may, for example, be connected to data processing systems 18.

10 Each end station 12 includes an authentication system 14. When an end station is sending data, the authentication system 14 included therein uses selected information and data bytes from one or more data packets to produce a corresponding set of one or more integrity checks. The authentication system next encrypts the integrity checks in a known manner in accordance with a shared secret key, and produces an “integrity block.” The
15 end station 12 then sends the data packets and the integrity block over the communications network 16 to one or more recipient end stations 12 through a network interface 15 (Fig. 2) in a known manner. The integrity block may be included in a data packet or it may be sent as a separate packet, as appropriate. The operations to produce the integrity checks and the integrity block are discussed below with reference to Figure 2. Hereinafter, we refer to the end stations 12 as senders and recipients based on their operations in
20 the transmission under discussion.

Referring now to Fig. 2, the authentication system 14 includes an integrity check processor 20 that, during send operations, produces the integrity checks. During receive operations the integrity check processor 20 uses the integrity checks to determine if the
25 received data packets are authentic, as discussed below. An integrity block processor 22, which is also included in the authentication system 14, encrypts the integrity checks and produces the integrity block that is sent over the network to the intended recipients. The integrity block processor 22 also decrypts a received integrity block, to reproduce the in-

egrity checks that are then used by the integrity check processor 20 to authenticate the received data packets.

We discuss below the basic authentication process performed by the authentication system 14 and thereafter several variations of that process. The basic authentication process uses shared secret keys. However, as discussed below, the sender and recipient
5 may also use a public/private key combination for further robustness.

The integrity check processor 20 selects from a first data packet “m” data bytes, with $1 \leq m \leq d$ and where “d” is the number of bytes in the data packet, and manipulates the bytes in accordance with a selected integrity check function to produce an associated
10 integrity check. The processor 20 in the example manipulates the selected bytes by concatenating them to produce the integrity check. The processor then supplies the integrity check to the integrity block processor 22 for encryption. The integrity check processor 20 similarly selects m corresponding bytes from the next “p”-1 data packets, with $p \geq 1$, manipulates the bytes on a per packet basis and supplies them to the integrity block processor 22 as the associated integrity checks.
15

When all p integrity checks are supplied to the integrity block processor 22, the processor encrypts them in a known manner using the shared secret key that is held in a register 24. The result of the encryption is an integrity block that is transferred over the communications network 16 to the intended recipients of the p data packets.

The integrity block processor 22 at the recipient readily decrypts the received integrity block in a known manner using the shared secret key, and reproduces the integrity checks for the p data packets. The integrity check processor 20 at the recipient then authenticates each of the p received data packets by manipulating selected bytes from the packet in accordance with the selected integrity check function, or in the example, compares the bytes in the associated integrity check with the corresponding bytes from the
20 appropriate data packet. If the bytes match, the integrity check processor determines that the given data packet is authentic. The authentication system 14 thus authenticates p data packets by performing a single decryption operation and p m-byte compare operations. This is in contrast with known prior systems that authenticate p data packets by performing
25 p separate authentication operations, each of which typically involves manipulating
30

all of the data bytes in the packet in accordance with a relatively complex encoding function. Accordingly, the system 14 authenticates the data packets faster than the known prior systems. Further, the encrypting of the integrity block allows the system to use a less robust, and thus, less complex process to produce the respective individual integrity checks. Accordingly, the system 14 may include less complex circuitry than prior known systems.

As discussed above, the authentication system may produce an encrypted integrity block for a single data packet. If the integrity block is included within the data packet, there will be no delay in authenticating the packet. Otherwise, the system must delay authentication of the packet until receipt of the associated integrity block.

If an interloper intercepts the data packets and the integrity block, it will not be able to decipher the encrypted integrity block, and thus, will not be able to determine which of the data bytes the recipient will use to test the authenticity of a given data packet and/or which integrity function to use. Accordingly, the interloper cannot readily determine from the integrity block how to alter a data packet and still satisfy the associated integrity check, even though the integrity check process is relatively weak. The integrity block, even if it relates only to a single data packet, is thus a relatively robust authentication tool because of the encryption of the individual integrity checks.

Further, the authentication process is not particularly complex or time consuming at the recipient end because the decryption operation is performed once for p data packets, with the decryption operation involving the relatively few bytes of the integrity checks. Also, the p integrity check operations are less complex since they may each involve a relatively small number of data bytes from each packet, and/or may use a relatively simple integrity check function that is easily and quickly performed, such as a CRC function because the selection of the integrity function is not known to unauthorized parties. Accordingly, the decryption and manipulation operations are less time consuming and less complex than the known prior similarly robust authentication techniques, which manipulate all of the data bytes of each of the p data packets in accordance with more complex functions to produce the integrity checks and later to test the authenticity of the p packets.

We discuss below further variations of the basic authentication process. Each variation adds more robustness to the process and, consequently, requires more time and adds complexity to the process at the recipient end.

As discussed above, the integrity check processor 20 at the sender selects m data
5 bytes from each data packet for the integrity checks. If corresponding bytes from each packet are used, an interloper who learns which particular bytes are selected can then alter one or more of the other bytes in a packet without detection. Similarly, if the interloper learns that only a relatively small number of bytes at, for example, the start of the packet, are selected for the integrity check, the interloper can then alter the data bytes at
10 the end of the packet without detection. Accordingly, the authentication system may instead select the data bytes at random from a first data packet and select from subsequent data packets bytes that are offset from those selected from the first packet. The system then includes in the encrypted integrity block information such as offset and interval values, and so forth, to inform the recipient which bytes to select for integrity checking from
15 each of the received data packets.

Alternatively, the system may select data bytes at random from each of the data packets, and include in the integrity block a random number seed value or other information that the recipient needs to reproduce the random byte selections for the various data packets.

20 The authentication system may also select from a known set of functions or family of functions the integrity check function to use for each packet. The system then includes in the integrity check, prior to encryption, a function identifier or some other information from which the recipient can determine the selected function.

As discussed, the integrity check may include one or more function identifiers
25 that specify to the recipient which integrity check function or functions to select for the associated data packet or packets. Alternatively, the sender and recipient may share a secret mechanism for determining a code to select the functions based on information in the integrity check or the corresponding data packet. For example, the sender and the recipient may share a secret seed that can be used to generate a pseudorandom sequence
30 through some function such as RC4. The recipient then uses the sequence number of the

packet or an offset value included in the integrity check as an offset into the sequence, and uses the next n bits to identify the selected identity check function.

The system may also select more than one integrity check function for use with a given data packet. The integrity check then includes a list of the identifiers or associated information and a corresponding list of the results of the encodings. The recipient may then choose one of the selected integrity check functions to authenticate the packet, and compare the result of the encoding of the received bytes using the chosen function with the received result that corresponds to the same function. To avoid detection, an interloper would thus have to produce a forged packet that decoded to the corresponding results for all of selected integrity functions, presumably without knowledge of which functions have been selected for the packet. Accordingly, the combining of more than one integrity functions for each of the p data packets further increases the robustness of the system over the use of a single function selected from the set of functions. This is true even if the set contains a relatively small number of functions.

As a further safeguard, the integrity check processor may encode the selected data bytes, that is, it may combine all or some of the selected bytes by, for example, XORing one or more of the bytes, performing ones complement addition with some or all of the bytes, selectively multiplying or adding the bytes, or performing similar logic functions, before the bytes are further manipulated using the selected integrity check functions. For even greater protection, the integrity check processor 20 may instead or in addition use a cryptographic hash function to encode the selected bytes or some combination of them before the bytes are encoded using the selected integrity check functions. An intermediate station that receives the packet may then determine that a given packet is valid using the selected integrity check function and the encrypted bytes, and leave for the intended destination station the full authentication of the packet using the cryptographic hash.

Before encrypting the integrity checks, the integrity block processor may arrange them in a different order than the associated data packets. At the same time, the system encrypts into the integrity block information that the recipient can use to reassign the integrity checks to the appropriate data packets. If, for example, the data packets include sequence numbers, the integrity block processor may include in the integrity block a list

of sequence numbers that associates the integrity checks with the data packets. Alternatively, each integrity check may include the sequence number of the associated data packet. The system may also include in a transmission an integrity block that contains the integrity checks for, for example, a next set of data packets or a previously sent set of data packets. The recipient then uses the information, such as the sequence numbers, included in the integrity block to determine how to assign the integrity checks reproduced from the decrypted integrity block to the appropriate data packets.

An interloper will presumably not be able to associate the out-of-order integrity checks with the appropriate data packets because of the encryption. Accordingly, the interloper will have an even harder task to determine how to alter the various data packets without detection.

The system may encrypt into the integrity block other information that a recipient requires to properly use the integrity checks. For example, the system may include in the integrity block executable code that implements a new integrity check process. The recipient then decrypts the integrity block to reproduce the individual integrity checks and the associated code, and runs the code to authenticate the data packets using the reproduced integrity checks. The system thus need not send the executable code separately to each recipient.

An interloper, even one who learns much about the old integrity check process, will not be able to decipher the code for the new process and will, therefore, not be able to determine how to alter the data bytes to avoid detection by the new integrity check process. For greater robustness, the sender may sign the code with its private key in a known manner, so that the recipient can be sure that an interloper has not altered the code.

Alternatively, the authentication system may include in the integrity block a digital signature that is produced in a known manner using the sender's private key. Before decrypting the integrity block, the recipient authenticates the integrity block using the sender's public key to ensure that the block has not been altered or forged.

For even further robustness, a sender may include in each data packet a digital signature that is encrypted with the sender's private key. A recipient can then decrypt a digital signature with the sender's public key, to double-check the authenticity of a packet that has survived the weaker integrity check process. The recipient can thus use the more
5 time consuming and complex, but stronger, authentication tool to spot check selected packets. With the spot check, the recipient can also detect data that has been forged by an interloper who has access to the shared secret key.

As discussed, the integrity blocks may be encrypted in a known manner, which typically involves providing to the recipients separate initialization vectors for each
10 block. The vectors are used in a known manner to set registers in the encryption hardware to a desired value before the hardware is used to encrypt, in this case, the integrity checks. To optimize the decryption process, the authentication system may instead use the system or method described in United States Patent No. 6,055,316 entitled System and Method for Deriving an Appropriate Initialization Vector for Secure Commu-
15 cations, which allows a recipient to derive the initialization vector directly from the received integrity block. The authentication system thus avoids having to send the vectors separately.

The authentication system may instead include in the integrity check prior to encryption a time stamp or sequence number. The time stamp or sequence number may
20 then be used as the initialization vector.

Referring now to Fig. 3, when the integrity block is produced for and included in a single data packet, the integrity block may be separated into two parts, namely, the identifiers for the selected integrity check functions and data bytes, and the results of manipulating the data bytes in accordance with the functions. The two parts may then be
25 separately encrypted so that the identifiers can be included at the start 32 of the data packet 30 and the results of the manipulation included at the end 34 of the data packet. The recipient then uses the function information at the start of the packet to set up the integrity check processor 20 (Fig. 2) and manipulates the data bytes as they are received, to produce a result that can be compared with the result included at the end of the received

packet. Before the two values are compared, however, either the result produced by the integrity check processor must be encrypted, or the received result decrypted.

Whether or not the integrity block is included in a data packet, the block may include in each integrity check one or more random bytes, to produce after encryption a value that is not readily decipherable without knowledge of the included random byte or bytes. Alternatively, the time stamp or sequence number discussed above may serve this purpose.

Generally, the interloper can not readily access the data packets. If this is not the case, the interloper may be able to capture one or more of them and produce packets that are only slightly altered, and may thus pass the integrity checks. To further inhibit the interloper, an alternative authentication system 14 depicted in Fig. 4 includes, in addition to the processors discussed above, a chaff processor 40 that introduces into the transmission one or more extraneous, or "chaff," data packets that intentionally fail the integrity checks. The chaff data packets have the same sequence numbers as the valid data packets, and an interloper without access to the information contained in the integrity block can not readily determine which of the same-numbered packets is the valid data packet. Accordingly, the interloper can not readily determine which packet to alter to produce a packet that may pass the integrity checks. The recipient, however, readily determines which of the packets are valid data packets, based on the integrity checks, and ignores the rest.

The chaff packets may be used with any of the authentication process variations discussed above. Namely, the chaff packets may be included in a transmission whether or not the integrity checks are concatenations of the selected data bytes or the results of the encoding of the selected data bytes, whether or not the integrity checks are assembled out of order, and so forth.

An interloper with access to the data packets may instead reproduce a valid data packet and send multiple copies of the packet in an attempt to monopolize network bandwidth. To avoid this problem, a recipient checks an included time stamp or packet sequence number against a clock or the values included in previous packets. The recipient then passes on or processes only those packets that have valid time stamps or num-

bers. For example, the recipient discards packets that include time stamps that fall behind the current time by more than a predetermined amount, where the amount is related to the time it takes the packet to traverse the network and the estimated clock skew among the stations. Similarly, the recipient discards packets with sequence numbers that fall too far behind the sequence numbers contained in the recently received packets.

The foregoing description has been limited to a specific embodiment of this invention. It will be apparent, however, that other processors and various forms of memory, including computer readable media, may be used for storing and executing program instructions pertaining to the techniques described herein. It will be further apparent that variations and modifications may be made to the invention, such as combining one or more of the processors listed separately into a single processor, including in the system stations that produce the encrypted integrity blocks but do not decrypt the integrity blocks and stations that decrypt the integrity blocks but do not encrypt them, and so forth, with the attainment of some or all of its advantages. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the invention.

What is claimed is: